RANDOM GROUPS CONTAIN SURFACE SUBGROUPS

DANNY CALEGARI AND ALDEN WALKER

ABSTRACT. A random group contains many quasiconvex surface subgroups.

1. INTRODUCTION

Gromov, motivated perhaps by the Virtual Haken Conjecture in 3-manifold topology (now a theorem of Agol [1]), famously asked the following

Surface Subgroup Question. *Let G be a one-ended hyperbolic group. Does G contain a subgroup isomorphic to the fundamental group of a closed surface with $\chi < 0$?*

Beyond its intrinsic appeal, one reason Gromov was interested in this question was a belief that such surface subgroups could be used as essential structural components of hyperbolic groups [9]. Our interest in this question is stimulated by a belief that surface groups (not necessarily closed, and possibly with boundary) can act as a sort of "bridge" between hyperbolic geometry and symplectic geometry (through their connection to causal structures, quasimorphisms, stable commutator length, etc.).

Despite receiving considerable attention the Surface Subgroup Question is wide open in general, although in the specific case of hyperbolic 3-manifold groups it was positively resolved by Kahn–Markovic [10]. Their argument depends crucially on the structure of such groups as lattices in the semisimple Lie group $\mathrm{PSL}(2,\mathbb{C})$. By contrast, in this paper we are concerned with much more combinatorial classes of hyperbolic groups. Nevertheless, one common point of contact between our approach and the approach of Kahn–Markovic is the use of probability theory to *construct* surfaces, and the use of (hyperbolic) geometry to *certify* them as injective. In particular, because our surfaces are certified as injective by local methods, they end up being quasiconvex. It is an interesting question to identify the class of hyperbolic groups which contain non-quasiconvex (yet injective) surface subgroups (hyperbolic 3-manifold groups are now known to contain such groups since they are virtually fibered, again by Agol [1]).

In [8], § 9 (also see [13]), Gromov introduced the notion of a *random group*, i.e. a group with fixed generating set and relations chosen randomly from the set of all (cyclically reduced) words of some length n with a fixed logarithmic density D. Properties of these groups are then shown to hold with probability going to 1 as $n \to \infty$; informally one says that a random group at density D has a certain property *with overwhelming probability*. Gromov showed that at density $D > 1/2$ such groups are trivial or $\mathbb{Z}/2\mathbb{Z}$, whereas at density $D < 1/2$ they are infinite, hyperbolic,

Date: April 9, 2013.

and two-dimensional (with overwhelming probability). Later, Dahmani–Guirardel–Przytycki [7] showed that they are one-ended and do not split, and therefore (by the classification of boundaries of hyperbolic 2-dimensional groups), have a Menger sponge as a boundary.

Random groups at density $D < 1/6$ are known to be cubulated (i.e. are equal to the fundamental groups of nonpositively curved compact cube complexes), and at density $D < 1/5$ to act cocompactly (but not necessarily properly) on a CAT(0) cube complex, by Ollivier–Wise [14]. On the other hand, groups at density $1/3 < D < 1/2$ have property (T), by Zuk [16] (further clarified by Kotowski–Kotowski [11]), and therefore cannot act on a CAT(0) cube complex without a global fixed point. A one-ended hyperbolic cubulated group contains a one-ended graph of free groups (see [6], Appendix A; this depends on work of Agol [1]), and Calegari–Wilton [6] show that a *random* graph of free groups (i.e. a graph of free groups with random homomorphisms from edge groups to vertex groups) contains a surface subgroup. Thus one might hope that a random group at density $D < 1/6$ should contain a graph of free groups that is "random enough" so that the main theorem of [6] can be applied, and one can conclude that there is a surface subgroup.

Though suggestive, there does not appear to be an easy strategy to flesh out this idea. Nevertheless in this paper we are able to show directly that at any density $D < 1/2$ a random group contains a surface subgroup (in fact, many surface subgroups).

We give three proofs of this theorem, valid at different densities, with the final proof giving any density $D < 1/2$. Theorem 5.2.4 is valid for one-relator groups (informally $D = 0$), Lemma 6.2.2 gives $D < 2/7$, while our main Theorem 6.4.1 gives $D < 1/2$. Explicitly, we show:

Surfaces in Random Groups 6.4.1. *A random group of length n and density $D < 1/2$ contains a surface subgroup with probability $1 - O(e^{-n^c})$. In fact, it contains $O(e^{n^c})$ surfaces of genus $O(n)$. Moreover, these surfaces are quasiconvex.*

This state of affairs is summarized in Figure 1. A modification of the construction (see Remark 6.4.2) shows that the surface subgroups can be taken to be homologically essential.

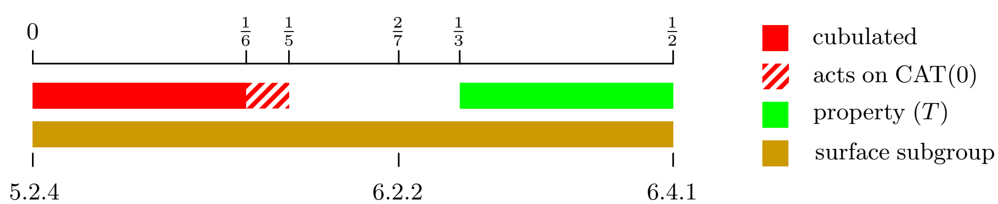

FIGURE 1. Random groups at different densities

Along the way we prove some results of independent interest. The first of these (and the most technically involved part of the paper) is the *Thin Fatgraph Theorem*, which says that a "sufficiently random" homologically trivial collection of cyclic words in a free group satisfies a strong combinatorial property: it can be realized as the boundary of a trivalent fatgraph in which *every* edge is longer than some prescribed constant:

Thin Fatgraph Theorem 3.3.1. *For all $L > 0$, for any $T \gg L$ and any $0 < \epsilon \ll 1/T$, there is an N depending only on L so that if Γ is a homologically trivial collection of tagged loops such that for each loop γ in Γ:*

(1) *no two tags in γ are closer than $4L$;*
(2) *the density of the tags in γ is at most ϵ/L;*
(3) *γ is (T, ϵ)-pseudorandom;*

then there exists a trivalent fatgraph Y with every edge of length at least L so that $\partial S(Y)$ is equal to N disjoint copies of Γ.

If the rank of the group is 2, we can take $N = 1$ above; otherwise we can take $N = 20L$.

The Thin Fatgraph Theorem strengthens one of the main technical theorems underpinning [5] and [6], and can be thought of as a kind of L^∞ theorem whose L^1 version (with optimal constants) is the main theorem of [3]. If r is a long random relator, the Thin Fatgraph Theorem lets us build a surface whose boundary consists of a small number of copies of r and r^{-1}. By plugging in a disk along each boundary component, we obtain a closed surface in the one-relator group $\langle F_k \mid r\rangle$. If the surface is built correctly, it can be shown to be π_1-injective, with high probability. This is one of the most subtle parts of the construction, and ensuring that the surfaces we build are π_1-injective at this step depends on the existence of a so-called *Bead Decomposition* for r; see Lemma 5.2.2. Thus we obtain the Random One Relator Theorem, whose statement is as follows:

Random One Relator Theorem 5.2.4. *Fix a free group F_k and let r be a random cyclically reduced word of length n. Then $G := \langle F_k \mid r\rangle$ contains a quasiconvex surface subgroup with probability $1 - O(e^{-n^c})$.*

The surfaces stay injective as more and more relators are added (in fact, these are the surfaces referred to in the main theorem) so this shows that random groups in the *few relators* model also contain surface subgroups, with high probability.

There is an interesting tension here: the fewer relators, the harder it is to build a surface group, but the easier it is to show that it is injective. This suggests looking for surface subgroups in an arbitrary one-ended hyperbolic group at a very specific "intermediate" scale, perhaps at the scale $O(\delta)$ where δ is the constant of hyperbolicity with respect to an "efficient" (e.g. Dehn) presentation.

1.1. **Acknowledgments.** We would like to thank Misha Gromov, John Mackay, Yann Ollivier, Piotr Przytycki and Henry Wilton. We also would like to acknowledge the use of Colin Rourke's `pinlabel` program, and Nathan Dunfield's `labelpin` program to help add the (numerous!) labels to the figures. Danny Calegari was supported by NSF grant DMS 1005246, and Alden Walker was supported by NSF grant DMS 1203888.

2. Background

In this section we describe some of the standard combinatorial language that we use in the remainder of the paper. Most important is the notion of *foldedness* for a map between graphs, as developed by Stallings [15]. We also recall some standard elements of the theory of small cancellation, which it is convenient to cite at certain points in our argument, though ultimately we depend on a more flexible version of

small cancellation theory developed by Ollivier [12] specifically for application to random groups (his results are summarized in § 6.1).

2.1. **Fatgraphs and foldedness.**

Definition 2.1.1. Let X and Y be graphs. A map $f : Y \to X$ is *simplicial* if it takes edges (linearly) to edges. It is *folded* if it is locally injective.

A folded map between graphs is injective on π_1. The terminology of foldedness, and its first effective use as a tool in group theory, is due to Stallings [15].

Definition 2.1.2. A *fatgraph* is a graph Y together with a choice of cyclic order on the edges incident to each vertex. A fatgraph admits a canonical *fattening* to a surface $S(Y)$ in which it sits as a spine (so that $S(Y)$ deformation retracts to Y) in such a way that the cyclic order of edges coming from the fatgraph structure agrees with the cyclic order in which the edges appear in $S(Y)$. A *folded fatgraph over* X is a fatgraph Y together with a folded map $f : Y \to X$.

The case of most interest to us will be that X is a rose associated to a free generating set for a (finitely generated) free group F.

A folded fatgraph $f : Y \to X$ induces a π_1 injective map $S(Y) \to X$. The deformation retraction $S(Y) \to Y$ induces an immersion $\partial S(Y) \to Y$, and we may therefore think of $\partial S(Y)$ as a union of simplicial loops. Under f these loops map to immersed loops in X, corresponding to conjugacy classes in $\pi_1(X)$.

Conversely, given a homologically trivial collection of conjugacy classes Γ in $\pi_1(X)$ represented (uniquely) by immersed oriented loops in X, we may ask whether there is a folded fatgraph Y over X so that $\partial S(Y)$ represents Γ (by abuse of notation, we write $\partial S(Y) = \Gamma$). Informally, we say that such a Γ *bounds a folded fatgraph*.

2.2. **Small cancellation.**

Definition 2.2.1. Let G have a presentation

$$G := \langle x_1, \cdots, x_n \mid r_1, \cdots, r_s \rangle$$

where the r_j are cyclically reduced words in the generators $x_i^{\pm}$. A *piece* is a subword that appears in two different ways in the relations or their inverses. A presentation satisfies the condition $C'(\lambda)$ for some λ if every piece σ in some r_i satisfies $|\sigma|/|r_i| < \lambda$.

Remark 2.2.2. Some authors use the notation $C'(\lambda)$ to indicate the weaker inequality $|\sigma|/|r_i| \le \lambda$. This distinction will be irrelevant for us.

Associated to a presentation there is a connected 2-complex K with one vertex, one edge for each generator, and one disk for each relation. The 1-skeleton X for K is a rose for the free group on the generators. As is well-known, a group satisfying $C'(1/6)$ is hyperbolic, and the 2-complex K is aspherical (so that the group is of cohomological dimension at most 2).

Definition 2.2.3. Fix a group G with a presentation complex K and 1-skeleton X as above. A *surface over the presentation* is an oriented surface S with the structure of a cell complex together with a cellular map to K which is an isomorphism on each cell. The 1-skeleton Y of the CW complex structure on S inherits the structure of a fatgraph from S and its orientation, and this fatgraph comes together with a map to X. We say S has a *folded spine* if $Y \to X$ is a folded fatgraph.

If G is a small cancellation group, a surface with a folded spine can be certified as π_1 injective by the following combinatorial condition.

Definition 2.2.4. Let G be a group with a fixed presentation, and let S be an oriented surface over the presentation with a folded spine Y. We say S is *α-convex* (for some $\alpha > 0$) with respect to the presentation if for every immersed path γ in Y which is a piece in some relation $r_i^{\pm}$ with $|\gamma|/|r_i| \ge \alpha$, we actually have that γ is contained in $\partial S(Y)$ (i.e. it is in the boundary of a disk of S).

Lemma 2.2.5 (Injective surface)**.** *Let G be a group with a presentation satisfying $C'(1/6)$, and let S be an oriented surface over the presentation with a folded spine Y. If S is $1/2$-convex then it is injective. Moreover, if S is α-convex for any $\alpha < 1/2$ then it is quasiconvex.*

Proof. First we prove injectivity under the assumption that S is $1/2$-convex. Suppose not, so that there is some essential loop in $\pi_1(S)$ which is trivial in G. After a homotopy, we can assume this loop γ is *immersed* in Y. Since Y is folded, the image of γ in X is also immersed; i.e. it is represented by a cyclically reduced word in the generators. Since by hypothesis γ is trivial in G, there is a van Kampen diagram with γ as boundary. We may choose γ and a diagram for which the number of faces is minimal.

The $C'(1/6)$ condition implies that there is a face D in the diagram which has at least $1/2$ of its boundary as a connected segment on γ. Then the hypothesis implies that this segment is actually contained in the boundary of a disk D' of S. Since G is $C'(1/6)$ it follows that D' and D bound the same relator in the same way, and we can therefore push γ across D' to obtain a van Kampen diagram with fewer faces and with boundary an essential loop in S (homotopic to γ). But this contradicts the choice of van Kampen diagram, and this contradiction shows that no such essential loop exists; i.e. that $\pi_1(S) \to G$ is injective.

The proof of quasiconvexity is similar. If S is not quasiconvex, there is a van Kampen diagram on an annulus A, with one boundary component a geodesic in G and the other a geodesic γ in Y, in which the ratio of the lengths of the boundary components is as big as desired. Since G is $C'(1/6)$ either the annulus has uniform thickness (which implies that the ratio of lengths is bounded) or else some face has a connected subpath of length at least α of its boundary on γ, for any $\alpha < 1/2$. The argument is then as above. $\square$

3. Trivalent fatgraphs

The purpose of this section is to prove the Thin Fatgraph Theorem 3.3.1, which implies that a (homologically trivial) collection of random cyclically reduced words bounds a trivalent fatgraph with long edges (i.e. in which *every* edge is as long as desired).

For concreteness the theorem is stated not for random words but for (sufficiently) pseudorandom words, and does not therefore really involve any probability theory. However the (obvious) application in this paper is to random words, and words obtained from them by simple operations.

3.1. Partial fatgraphs and tags.

We are going to build folded fatgraphs with prescribed boundary (i.e. given Γ we will build Y with $\Gamma = \partial S(Y)$). In the process of building these fatgraphs we deal with intermediate objects that we call *partial*

fatgraphs bounding part of Γ, and the part of Γ that is not yet bounded by a partial fatgraph is a collection of *cyclic words with tags*. This language is introduced in [5].

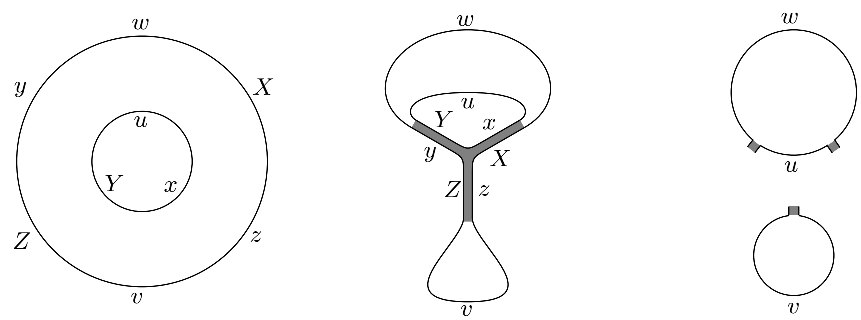

FIGURE 2. Two cyclic words are partially paired along a partial fatgraph (the grey tripod); what is left is two cyclic words with three tags.

The (partial) fatgraphs will be built by taking disjoint pairs of segments in Γ with inverse labels (in X) and *pairing them* — i.e. associating them to opposite sides of an edge of the fatgraph. Once all of Γ is decomposed into such paired segments the fatgraph Y will be implicitly defined.

A *partial fatgraph* is, abstractly, the data of a pairing of some collection of disjoint pairs of segments in Γ. We imagine that this partial fatgraph Z has boundary $\partial S(Z) =: \Gamma'$ which is a subset of Γ. The difference $\Gamma - \Gamma'$ is a collection of paths whose endpoints are paired according to how they are paired in Γ'. The result is therefore a collection of cyclic words Γ'', together with the data of the "germ" of the partial fatgraph Z at finitely many points. This extra data we refer to as *tags*, and we call this collection Γ'' a collection of *cyclic words with tags*. See Figure 2 for an example.

3.2. **Pseudorandomness.** Random (cyclically reduced) words enjoy many strong equidistribution properties, at a large range of scales. For our purposes it is sufficient to have "enough" equidistribution at a sufficiently large *fixed* scale. To quantify this we describe the condition of *pseudorandomness*, and observe that random words are pseudorandom with high probability.

Definition 3.2.1. Let Γ be a cyclically reduced cyclic word in a free group F_k with $k \geq 2$ generators. We say Γ is (T, ϵ)*-pseudorandom* if the following is true: if we pick any cyclic conjugate of Γ, and write it as a product of reduced words w_i of length T (and at most one word v of length $< T$)

$$\Gamma := w_1 w_2 w_3 \cdots w_N v$$

then for every reduced word σ of length T in F_k, there is an estimate

$$1 - \epsilon \leq \frac{\#\{i \text{ such that } w_i = \sigma\}}{N} \cdot (2k)(2k-1)^{T-1} \leq 1 + \epsilon$$

Similarly, we say that a collection of reduced words w_i of length T is ϵ*-pseudorandom* if for every reduced word σ of length T in F_k the estimate above holds.

Lemma 3.2.2 (Random is pseudorandom)**.** *Fix $T, \epsilon > 0$. Let Γ be a random cyclically reduced word of length n. Then Γ is (T, ϵ)-pseudorandom with probability $1 - O(e^{-Cn})$.*

Proof. This is immediate from the Chernoff inequality for finite Markov chains. $\square$

3.3. **Thin fatgraph theorem.** We now come to the main result in this section, the Thin Fatgraph theorem. This says that any (sufficiently) pseudorandom homologically trivial collection of tagged loops, with sufficiently few and well-spaced tags, bounds a trivalent fatgraph with every edge as long as desired. Note that every trivalent graph (with reduced boundary) is automatically folded.

This theorem can be compared with [5], Thm. 8.9 which says that random homologically trivial words bound 4-valent folded fatgraphs, with high probability; and [3], Thm. 4.1 which says that random homologically trivial words of length n bound (not necessarily folded) fatgraphs whose *average* valence is arbitrarily close to 3, and whose *average* edge length is as close to $\log(n)/2\log(2k-1)$ as desired (and moreover this quantity is sharp). It would be very interesting to prove (or disprove) that random homologically trivial words bound (with high probability) trivalent fatgraphs in which *every* edge has length $O(\log(n))$, but this seems to require new ideas.

Theorem 3.3.1 (Thin Fatgraph)**.** *For all $L > 0$, for any $T \gg L$ and any $0 < \epsilon \ll 1/T$, there is an N depending only on L so that if Γ is a homologically trivial collection of tagged loops such that for each loop γ in Γ:*

1. *no two tags in γ are closer than $4L$;*
2. *the density of the tags in γ is at most ϵ/L;*
3. *γ is (T, ϵ)-pseudorandom;*

then there exists a trivalent fatgraph Y with every edge of length at least L so that $\partial S(Y)$ is equal to N disjoint copies of Γ.

The notation $T \gg L$ means "for all T sufficiently long depending on L", and similarly $0 < \epsilon \ll 1/T$ means "for all ϵ sufficiently small depending on T". The role of N will become apparent at the last step, where some combinatorial condition can be solved more easily over the rationals than over the integers (so that one needs to take a multiple of the original chain in order to clear denominators). In fact, in rank 2 we can actually take $N = 1$, and in higher rank we can take $N = 20L$ (it is probably true that one can take $N = 1$ always, but this is superfluous for our purposes).

Except for the last step (which it must be admitted is quite substantial and takes up almost half the paper), the argument is very close to that in [5]. For the sake of completeness we reproduce that argument here, explaining how to modify it to control the edge lengths and valence of the fatgraph.

3.4. **Experimental results.** Theorem 3.3.1 asserts that long random words bound trivalent fatgraphs (up to taking sufficiently many disjoint copies). However, in order for the pseudorandomness to hold at scales required by the argument, it is necessary to consider random words of enormous length; i.e. on the order of a googol or more. On the other hand, experiments show that even words of modest length bound trivalent fatgraphs with high probability. To keep our experiment simple, we considered only the condition of bounding a trivalent graph, ignoring the question of whether the edges can all be chosen to be long.

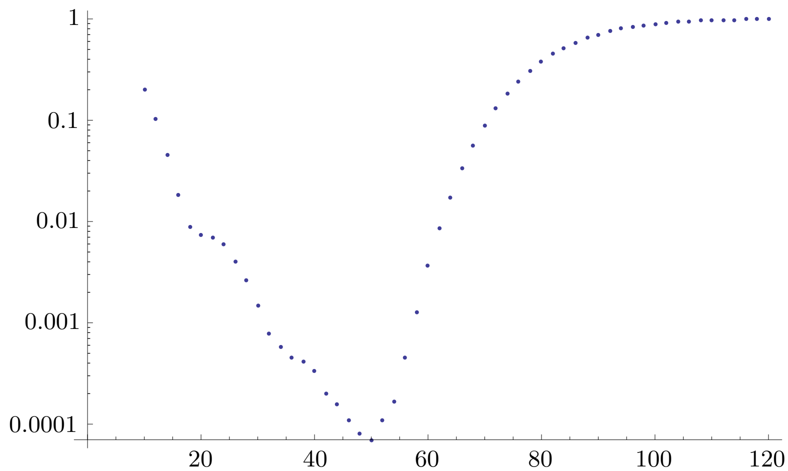

Figure 3. The experimental log-fraction of random words of each length which bound a trivalent fatgraph.

In a free group of rank 3, we looked at between 100000 and 400000 cyclically reduced homologically trivial words of each even length from 10 to 120. The proportion of such words that bound trivalent fatgraphs is plotted in Figure 3. The vertical axis has a log-scale to show some interesting features of the data. As one can see, bounding a trivalent fatgraph happens in practice for n far below the purview of Theorem 3.3.1.

3.5. **Proof of the thin fatgraph theorem.** We now give the proof of Theorem 3.3.1. The proof proceeds in several steps. The first few steps are more probabilistic in nature. The last step is more combinatorial and quite intricate, and is deferred to § 4.

Pick a γ in Γ. Choose any location in γ and write

$$\gamma = w_1 w_2 \cdots w_N v$$

where each $|w_i| = T$ and $N = \lfloor |\gamma|/T \rfloor$. Since γ is (T,ϵ)-pseudorandom, the w_i are very well equidistributed among the reduced words of length T. Moreover, since the density of tags is sufficiently low, very few w_i contain a tag. We restrict attention to the w_i that do not contain a tag.

3.5.1. *Tall poppies.* Throughout the remainder of the proof we fix some T' which is an odd multiple of $10L$ with $1000L < T - T' \le 2000L$. Note that we still have $T' \gg L$. For each w_i we let v_i be the initial subword of length T'. Note that the map which takes a reduced word of length T to its prefix of length T' takes the uniform measure to a multiple of the uniform measure, and therefore the v_i are also ϵ-pseudorandom.

The first step is to create a collection of *tall poppies*. We fix some $v := v_i$ and read the letters one by one. As we read along, we look for a pair of inverse subwords x, X each of length $10L$ and separated by a subsegment y of length $40L$. Further we require that the copy of xyX should have the property that the x and X are maximal inverse subwords at their given locations, so that the result of pairing creates reduced tagged cyclic words. We create some partial fatgraph by identifying

x to X; this creates a *tall poppy* whose *stem* is x, and whose *flower* is y. Once we find and create a tall poppy, we look for each subsequent tall poppy at successive locations along v subject to the constraint that adjacent tall poppies are separated by subwords whose length is an even multiple of $10L$. Furthermore, we insist that the first tall poppy occurs at distance an even multiple of $10L$ from the start of v. See Figure 4 for an example; the "dots" in the figure indicate units of $10L$.

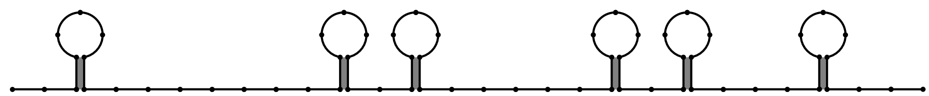

FIGURE 4. A word v of length $630L$ with 6 tall poppies folded off

For each v_i we fold off tall poppies as above. The result of this step is to create a partial fatgraph for each γ consisting of some tagged loop γ' (which is obtained from γ by cutting out all the xyX subwords and identifying endpoints) and a *reservoir* of flowers. Observe that every tagged cyclic word of length $40L$ occurs as a flower, and the set of tagged flowers is ϵ-pseudorandom (conditioned on any compatible label on the tag). Informally, we say that the reservoir contains an *almost equidistributed collection* of tagged cyclic words of length $40L$. We can estimate the total number of flowers of each kind: at each location that a flower might occur, we require two subwords of length $10L$ to be inverse, which will happen with probability $(2k-1)^{-10L}$. The number of locations is roughly of size $O(|\gamma|/10L)$. So the number of copies of each tagged loop in the reservoir is of size $\delta \cdot |\gamma|$ (up to multiplicative error $1 \pm \epsilon$) for some specific positive $\delta > 0$ depending only on L.

3.5.2. *Random cancellation.* After cutting off tall poppies, the v_i become tagged words v_i'. Observe that the v_i' have variable lengths (differing from T' by an even multiple of $10L$) and have tags occurring at some subset of the points an even multiple of $10L$ from the start. The main observation to make is that the ϵ-pseudorandomness of the v_i propagates to ϵ-pseudorandomness of the v_i'. That is, if σ is a reduced word of length $T' - m10L > 0$ for some even m, then among the v_i' of length $T' - m10L$, the proportion that are equal to σ is equal to $1/(2k)(2k-1)^{T'-m10L-1}$ up to a multiplicative error of size $1 \pm \epsilon$. This is immediate from the construction.

Recall that we chose T' to be an *odd* multiple of $10L$. This means that when we pair a segment v_i' labeled σ with a v_j' labeled σ^{-1} the tags of v_i' and v_j' do not match up, and in fact any two tags are no closer than distance L. In fact, it is important that after pairing up inverse segments, the tagged loops that remain are reduced, so we write each v_i' in the form $l_i v_i'' r_i$ where each of l_i, r_i has length $5L$, and pair v_i'' with v_j'' for some v_j' of the form $l_j v_j'' r_j$ where $v_j'' = (v_i'')^{-1}$, and the words $l_i r_j$ and $r_i l_j$ are reduced. By ϵ-pseudorandomness, we can find such pairings of all but $O(\epsilon)$ of the v_i' in this way.

Thus the result of this pairing is to produce a trivalent partial fatgraph with all edges of length at least $10L$. Removing this from the v_i' produces a collection of tagged loops γ'' with $|\gamma''| = O(\epsilon \cdot |\gamma|)$.

3.5.3. *Cancelling Γ'' from the reservoir.* Let Γ'' be the union of all the γ'', and pool the reservoirs from each γ into a single reservoir.

Notice that by construction, no tagged loop in Γ'' has two tags closer than distance $20L$. It is clear that for each tagged loop ν in Γ'' we can build a copy of ν^{-1} out of finitely many flowers in the reservoir in such a way that the result of pairing ν to this ν^{-1} is a trivalent partial fatgraph with all edges of length at least L, and the number of flowers that we need is proportional to $2|\nu|/L$. There is a slight subtlety here, in that the length of each flower is $40L$, and the result of partially gluing up a collection of cyclic words of even length always leaves an even number of letters unglued. Fortunately, the assumption that Γ is homologically trivial implies that $|\Gamma|$ itself is even, and since each flower also has an even number of letters, it follows that $|\Gamma''|$ is even. A flower with the cyclic word $xyzY$ can be partially glued to produce two tagged loops x and z, and if x and z are odd, each can be used to contribute to a copy of some ν^{-1} of odd total length. Since the number of odd $|\nu|$ is even, all of Γ'' can be cancelled in this way.

Since $|\Gamma''| = O(\epsilon\cdot|\Gamma|)$ whereas the number of flowers of each kind in the reservoir is of order $\delta\cdot|\Gamma|$, if we take $\epsilon \ll \delta$ we can glue up all of Γ'' this way, at the cost of slightly adjusting the proportion of each kind of tagged loop in the reservoir.

3.5.4. *Gluing up the reservoir.* We are now left with an almost equidistributed collection of tagged loops of length $40L$ in the reservoir. Adding to the reservoir the contribution from each γ in Γ, and using the fact that Γ was homologically trivial, we see that the content of the reservoir is also homologically trivial. It remains to show that any such collection can be glued up to build a trivalent partial fatgraph with all edges of length at least L.

In fact, we only need two kinds of gluings to achieve this: gluings that result in partial fatgraphs that fatten to *annuli* and to *pants*. The argument is purely combinatorial, but quite intricate and involved, and takes up the content of § 4.

4. Annulus moves and pants moves

In this section we show that an almost equidistributed collection of tagged loops of length $40L$ can be glued up to a trivalent partial fatgraph with all edges of length at least L. Together with the content of § 3.5, this will conclude the proof of Theorem 3.3.1.

Remark 4.0.1. For the entirety of this section, we will rescale $40L$ to L. That is, we prove that an almost equidistributed collection of tagged loops of length L, where L is divisible by 4, can be glued up to a trivalent partial fatgraph with all edges of length at least $L/4$. This rescaling is intended to remove meaningless factors of 40 throughout the argument.

4.1. **Pants and annuli.** Let $S(L)$ be the set of tagged loops of length L, where L is divisible by 4. Let $W(L)$ be the vector space over $\mathbb{Q}$ spanned by $S(L)$; that is, $W(L) = \mathbb{Q}[S(L)]$. We define $h : W(L) \to \mathbb{Z}\times\mathbb{Z}$ to be the linear map so that $h(v)$ is the homology class of v. Finally, $V(L) = \ker h \subseteq W(L)$ is the vector space of homological trivial vectors. We are interested only in $V(L)$, not $W(L)$, so by "full dimensional", we mean a full dimensional subset of $V(L)$. When we say that a vector *projectively bounds* a fatgraph, we mean that there is some multiple of the vector which has integer coordinates, and the collection of loops represented by the integral vector bounds a fatgraph. A (necessarily integral) vector *bounds* a fatgraph

if the collection of loops that it represents bounds a fatgraph. The uniform vector of all 1's will be of particular interest, and we denote it by $\mathbf{1}$.

We say that a fatgraph Y with boundary a collection of loops in $S(L)$ is *thin* if Y is trivalent and the trivalent vertices of Y are pairwise distance at least $L/4$ apart, where the tags are counted as trivalent vertices. Let $C(L)$ be the subset of $V(L)$ of positive vectors which projectively bound a thin fatgraph. If $v, w \in C(L)$, then the disjoint union of the thin fatgraphs for v and w gives a thin fatgraph for $v+w$. Also, the definition of $C(L)$ shows it to be closed under scalar multiplication. Hence, $C(L)$ is a cone. A variant of the `scallop` [4] algorithm gives an explicit hyperplane description of $C(L)$, and shows that it is a finite sided polyhedral cone, but we won't need this fact in the sequel.

We will build thin fatgraphs out of two kinds of pieces: (good pairs of) *pants* and (good) *annuli* (the terminology is supposed to suggest an affinity with the Kahn–Markovic proof of the Ehrenpreis conjecture, but one should not make too much of this). A good pair of pants is one whose edge lengths are all exactly $L/2$ and whose tags are each on different edges and exactly distance $L/4$ from the real trivalent vertices. Note the boundary of each such pair of pants lies in $V(L)$. A good annulus is a fatgraph annulus with boundary in $S(L)$ whose tags are distance at least $L/4$ apart. Hereafter, all pants and annuli are good.

Define an involution $\iota : S(L) \to S(L)$ which takes each loop to its inverse with the tag moved to the diametrically opposite position. There are several options for the tag at each position – for the definition of ι, we arbitrarily choose any pairing of the options to obtain an involution. There is a special class of annuli, which we call ι-annuli, which have boundary of the form $s + \iota(s)$. Notice that the collection of all ι-annuli is a thin fatgraph which bounds the uniform vector $\mathbf{1}$.

The bulk of our upcoming work lies in manipulating *un*tagged loops, and our result here is independently interesting, so we will need some complementary definitions. Let $S'(L)$ be the set of untagged loops of length L, let $W'(L) = \mathbb{Q}[S'(L)]$, and let $V'(L)$ be the vector space of homologically trivial vectors in $W'(L)$. We define a thin fatgraph and the uniform vector $\mathbf{1}' \in V'(L)$ as before. The set $C'(L) \subseteq V'(L)$ is the cone of vectors in $V'(L)$ which projectively bound thin fatgraphs. An untagged good pair of pants is a trivalent pair of pants whose edge lengths are exactly $L/2$, and an untagged annulus is simply an annulus whose boundary is two loops of length L. For untagged loops, $\iota : S' \to S'$ is simply inversion, and all annuli are ι-annuli, although we may refer to them explicitly as ι-annuli to emphasize their purpose.

For many applications, the property of a collection of loops that it *projectively* bounds a thin fatgraph is good enough (see e.g. [5]), and this is in many ways a more pleasant property to work with, since the set of vectors (representing collections of loops) which projectively bound a thin fatgraph is a cone, whereas the set of vectors that *bound* (i.e. without resorting to taking multiples) is the intersection of this cone with an integer lattice. However, in this paper it is important to distinguish between "bounding" and "projectively bounding", and therefore in the following propositions, we give both the stronger, technical "integral" statement and the weaker, cleaner "rational" one.

Proposition 4.1.1. *For any integral vector $v \in V'(L)$, there is $n \in \mathbb{N}$ so that $(L/2)v + n\mathbf{1}'$ bounds a collection of good pants and annuli. Consequently, $C'(L)$ is full dimensional and contains an open projective neighborhood of $\mathbf{1}'$.*

There is a stronger version without the $L/2$ factor if the free group has rank 2.

Proposition 4.1.2. *If the free group has rank 2, then for any integral vector $v \in V'(L)$, there is $n \in \mathbb{N}$ so that $v + n\mathbf{1}'$ bounds a collection of good pants and annuli.*

We believe that Proposition 4.1.2 is probably true for higher rank, but the proof would be more complicated than we wish for a detail that we do not need.

We delay the rather tedious proof of Proposition 4.1.1 in favor of stating the tagged version, which is a corollary and is the version we need.

Proposition 4.1.3. *For any integral vector $v \in V(L)$, there is $n \in \mathbb{N}$ so that $(L/2)v + n\mathbf{1}$ bounds a collection of good pants and annuli. Consequently, $C(L)$ is full dimensional and contains an open projective neighborhood of $\mathbf{1}$.*

Proof. Let us be given an integral vector $v \in V(L)$. Define $f : V(L) \to V'(L)$ to be the map $S(L) \to S'(L)$ which forgets the tag, extended by linearity. By Proposition 4.1.1, we can find a collection of pants and annuli which has boundary $(L/2)f(v) + m\mathbf{1}'$. Call this fatgraph Y'. Now place arbitrary tags in the forced positions on the pants (in the middle of the edges) and in allowed positions on the annuli (at least $L/4$ apart) to obtain Y. Clearly, Y is a thin fatgraph, and Y almost has boundary $(L/2)v + m\mathbf{1}$, as desired, but the tags are in the wrong places. We fix this by simply adding annuli which "twist" the tags into the right positions.

For each boundary of Y corresponding to a loop γ in $(L/2)v$, we can add at most two annuli to Y to resolve the boundary into $\gamma + \alpha + \iota(\alpha)$, where α is one or two tagged loops; see Figure 5. The remaining boundary of Y consists of a collection of inverse pairs in $V'(L)$, but is now incorrectly tagged, and so does not decompose into ι pairs. Choose a "correct" desired tag arbitrarily for each boundary pair so that they constitute an ι orbit γ and $\iota(\gamma)$. In the same way as we did for the loops in $(L/2)v$, we can add (at most 4) annuli to resolve the boundary into the correct $\gamma + \iota(\gamma) + \alpha + \iota(\alpha)$. The boundary of Y now decomposes into $(L/2)v + w + \iota(w)$ for some integral $w \in V(L)$. By adding ι-annuli, we can make the boundary of Y be $(L/2)v + n\mathbf{1}$, for some $n \geq m$, as desired.

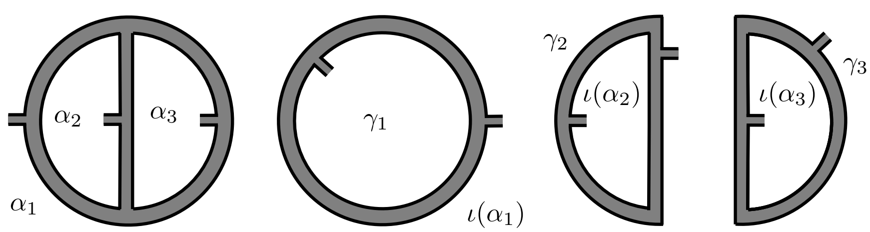

FIGURE 5. Adding annuli to fix incorrectly tagged boundary. We desire the tagged boundary $\gamma_1 + \gamma_2 + \gamma_3$. Applying Prop. 4.1.1, we find the pair of pants, left, with the correct boundary, but tags in the wrong places. Adding a few annuli produces the correct, tagged boundary, plus a collection of ι orbits $\alpha_i + \iota(\alpha_i)$.

∎

It remains to prove Proposition 4.1.1, which we now do.

4.2. **Proof of Proposition 4.1.1.** Let us be given an integral vector $v' \in V'(L)$. The vector v' represents a collection of tagged loops $s' \in S'(L)$, for which we must find a thin fatgraph Y of the desired form. Our goal is to build a fatgraph which bounds $s' + t + \iota(t)$ for some $t \in S'(L)$. Adding ι-annuli will then immediately finish the construction.

If we have a collection of annuli and pants which has boundary $s' + t$, then the problem reduces to finding a collection of annuli and pants which has boundary of the form $\iota(t) + u + \iota(u)$ for some $u \in S'(L)$. We'll repeatedly apply this idea to simplify the problem by *attaching pants*. If we want to have boundary which contains a loop γ, and we find a pair of pants with boundary $\gamma + \alpha + \alpha'$, then now we need only find boundary containing $\iota(\alpha) + \iota(\alpha')$. In this case, we'll say that γ and $\iota(\alpha) + \iota(\alpha')$ are *pants equivalent*.

For this entire section, we will assume that our free group has rank 2 and is generated by a and b. In § 4.3 we explain the extra details required to deal with free groups of higher rank.

Our strategy will be to start with s' and attach (many) pairs of pants which put all the loops in s' into a nice form. Then we attach more pants to further simplify the loops, and so on, eventually reducing to a case that is simple enough to handle by hand. A *run* in a loop is a maximal subword of the form a^p or b^p for some integer power $p \neq 0$. Note that any loop contains an even number of runs. A loop has *balanced a runs*, respectively b runs, if all the runs of a's, respectively b's, have the same length. A loop has *balanced runs* if it has balanced a runs and balanced b runs. We first show how to attach pants to reduce any loop to a collection of loops with at most 4 runs. Then we attach pants to produce loops with balanced runs, then loops with 2 runs, then we simplify this collection into a standard form.

For clarity, we separate the simplification into lemmas.

Lemma 4.2.1. *Any loop is pants-equivalent to a loop with* 4 *runs.*

Proof. To begin, we show how to attach a pair of pants to a loop which produces two loops, each of which has fewer runs than the initial loop. This method works whenever the number of runs is more than 4, so it reduces the loops in s' to a collection of loops with at most 4 runs. Let us be given a loop γ. The easiest way to visualize attaching a pair of pants is to simply draw a diameter d on γ between two antipodal vertices. Labeling the diameter d produces a pair of pants attached to γ. Note that we must be careful to label d compatibly with the labels adjacent to the vertices to which we attach d, so that the vertices do not fold. See Figure 6.

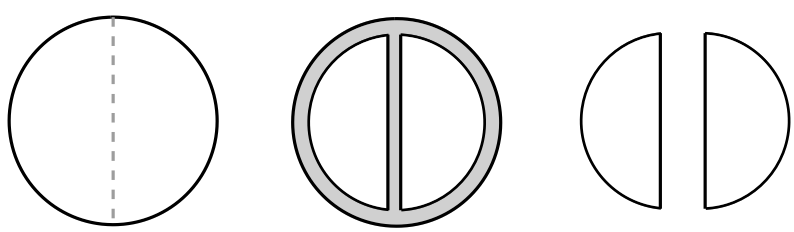

FIGURE 6. Attaching a diameter to a loop to form a pair of pants. The pair of loops on the right is pants equivalent to the loop on the left.

For concreteness, let us number the vertices of γ by $0, \ldots, L-1$, and we let d_i be the (oriented) diameter with initial vertex i (and thus terminal vertex $(i + L/2) \bmod L$). Each diameter d_i divides γ into two pieces. Let x_i be the number of runs in the non-cyclic subword of γ starting at index i and of length $L/2$; that is, the number of runs in the word to the "right" of d_i. Similarly, let y_i be the number of runs to the left. See Figure 7. Let r be the number of runs in γ. Note that $x_i + y_i$ may be greater than r. Specifically, $r \le x_i + y_i \le r + 2$. The important feature of these numbers is that $|x_i - x_{i+1}| \le 1$ and $|y_i - y_{i+1}| \le 1$. This is easily seen by considering the combinatorial possibilities that occur as we rotate the starting point i around the loop γ. See Figure 7. We are particularly interested in *matched runs*, which are runs separated in either direction by the same number of other runs. That is, matched runs are "directly across" from one another in the list of runs (we use scare quotes to emphasize that matched runs are not antipodal in the same sense that "antipodal vertices" are).

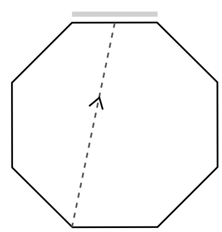

FIGURE 7. Moving the diameter one position changes x_i and y_i by at most 1. For the marked diameter, we have $x_i = 5$ and $y_i = 4$. A pair of matched runs is marked in grey.

If we connect the dots on the graphs of the functions x_i and y_i, then by the intermediate value theorem, there is some point at which the graphs of x_i and y_i intersect. This can happen either at some value of i, in which case $x_i = y_i$, or between two values i and $i+1$, but by the above fact, in this latter case $|x_i - y_i| \le 1$ and $|x_{i+1} - y_{i+1}| \le 1$. In either case, d_i must intersect two matched runs R and R', perhaps on the boundaries of the runs. Now decrease i until one of the ends of d_i lies on the boundary of R or R'. This puts d_i in to one of two combinatorial configurations, up to rotation and symmetry. See Figure 8. Note that the configuration on the right cannot occur at the intersection point, since $|x_i - y_i| = 2$.

We handle the two cases separately. First, the more generic case illustrated in Figure 8 on the left. Here we label the diameter entirely with the generator which is *not* the one labeling the bottom run, and in such a way as to minimize the number of resulting runs. Figure 9, left, illustrates this. The sign of the labels on the diameter depends on the signs and orders of the generators around the endpoints, but the picture is equivalent. Notice that the number of runs in each of the resulting loops is at most $r/2 + 2$.

In the non-generic case illustrated Figure 8 in the middle, we label half of the diameter with one generator and the other half with the other, in a way which is compatible with the top and bottom labels. See Figure 9, right. In certain cases, it is possible to label the entire diameter with a single generator, and this reduces

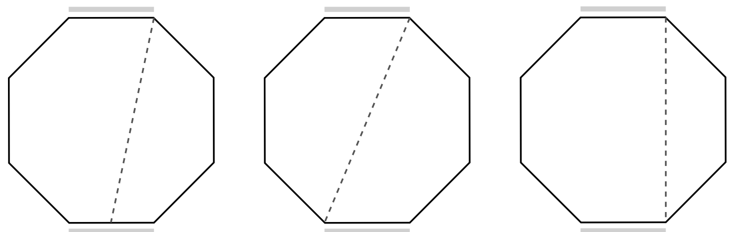

FIGURE 8. Possible configurations of d_i with respect to the matched runs R and R'. Up to rotation and symmetry, there are two. Note the configuration on the right cannot occur.

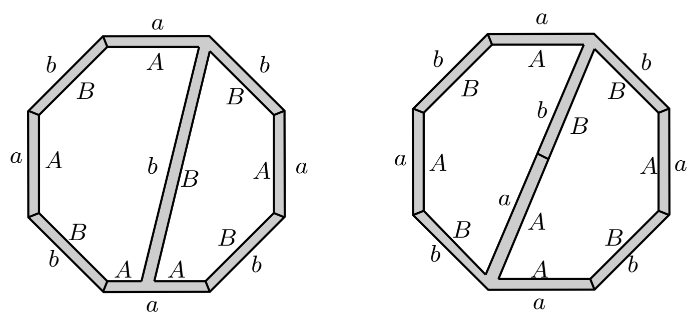

FIGURE 9. Labeling the diameter to reduce the number of runs. Each label represents a potentially long run of that generator.

the number of runs still further, but we have illustrated the worst situation. We therefore compute again that the number of runs in each of the resulting loops is $r/2 + 2$.

As long as $r/2 + 2 < r$, or $r > 4$, this will produce two loops of strictly smaller length. Repeatedly attaching pants resolves our collection s' into a new collection of loops, all of which have at most 4 runs. $\square$

Lemma 4.2.2. *Any loop with* 4 *runs is pants equivalent to a collection of loops with at most* 2 *runs.*

Proof. This is the most tedious section of the proof, and it relies on balancing out the 4 runs. There are two basic moves that we will use, a triangle move, and a division move. These are illustrated in Figure 10. The division move is self-explanatory. A correct way to think of the triangle move is to imagine rotating around the labels on the cut corner to label the new diameter.

Algebraically, a triangle move takes in a word $a^{e_1}b^{e_2}a^{e_3}b^{e_4}$ (for clarity, assume all the exponents are positive) such that $e_2 + e_3 > L/2$ and $e_2 < L/2$ and builds a pair of pants with boundary

$$a^{e_1}b^{e_2}a^{e_3}b^{e_4} + A^{e_1+L/2-e_2}B^{e_4}A^{e_3-(L/2-e_2)}B^{e_2} + a^{L/2-e_2}b^{e_2}A^{L/2-e_2}B^{e_2}$$

The notation obscures the function of a triangle move, which is as follows: it produces one loop which is balanced, and another loop with the same run sizes as the original one, except that $L/2 - e_2$ of the a's in the top run have been shifted down to the bottom run. The signs of the generators may change as they shift, but

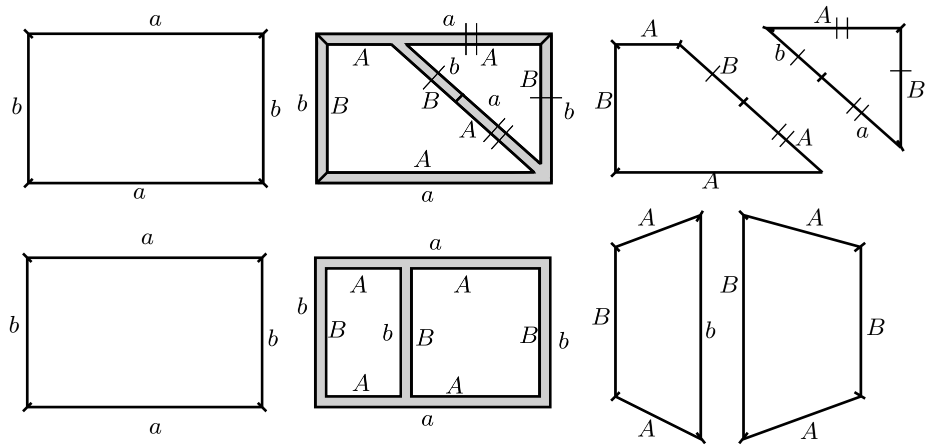

FIGURE 10. Using triangle and division moves to rearrange the runs in a loop. Lengths aren't to scale, except where marked to indicate runs of the same length

we are only concerned with the lengths at this point. This is the critical feature of the triangle moves — shifting generators from one run to the other without disturbing anything else.

Our goal will be to rearrange our 4-run loop into a collection of loops, each with a diameter *between corners*; that is, a diameter which begins at the junction of two runs and ends at the junction of the other two runs. Given a loop with a diameter between corners, there are two possibilities. Either no inverse pairs of generators appear, or there is at least one inverse pair. In the former case, a single triangle move produces a 2-run loop, and a loop with an inverse pair. In the latter, we can label the diameter with a single generator to produce two 2-run loops. Incidentally, performing this operation on a commutator produces two inverse loops. The two possibilities are shown In Figure 11.

Of course, not every 4-run loop has a diameter between corners. However, there is a sequence of moves that produces a collection of loops with diameters between corners. There are three possible inputs to this sequence, illustrated in Figure 12:

(1) an *equidistributed* loop;
(2) a *half-equidistributed* loop; or
(3) a loop which has two adjacent runs of lengths $L/4$ and $L/2$, respectively.

An equidistributed loop is a loop which has two diameters at right angles (with endpoints equally spaced) such that one diameter begins on one b run and ends on the other, and the other diameter begins on one a run and ends on the other. It is acceptable for the diameters to begin or end on a vertex between runs. A half-equidistributed loop is a loop for which one of the runs has length $L/2$, and both runs of the other generator have length at most $L/4$.

The sequence of moves which takes any of these three types of loops and produces loops with diameters between corners is shown in Figure 13. The output, at bottom right, is a perfectly balanced loop, which trivially has a diameter between corners, and the loop with the diameter shown. Although it is not to scale, this latter loop does have a diameter between corners, since it is the product of a triangle move on the loop to the left, and thus has two adjacent runs, both of length $L/4$.

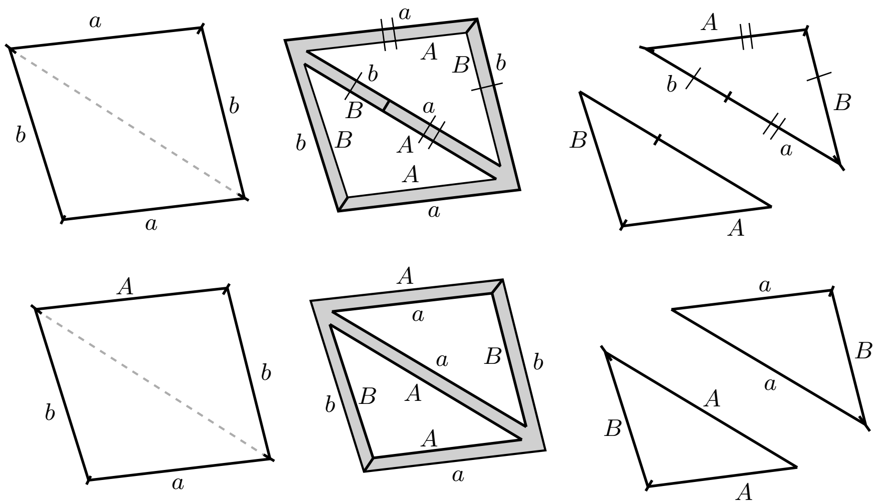

FIGURE 11. Performing triangle moves on loops with diameters between corners in order to produce 2-run loops. There are two possibilities, depending on whether an inverse pair appears.

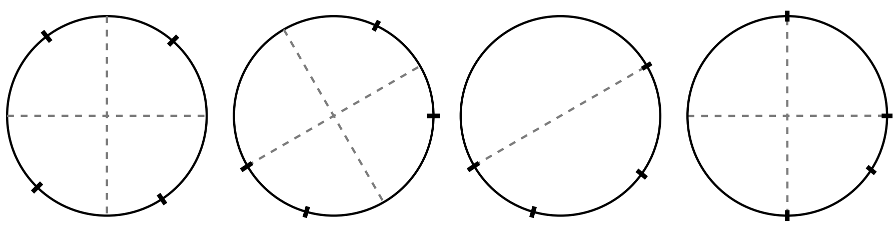

FIGURE 12. Two loops which are both equidistributed (left), a loop which is half-equidistributed, and a loop with adjacent runs of lengths $L/2$ and $L/4$.

So producing 2-run loops reduces to showing that we can make any loop equidistributed or half-equidistributed. For this, we'll use triangle moves to shift runs to appropriate locations.

Before moving on to the next step of the argument, we remark that the triangle and division moves can be reinterpreted geometrically in the following way. Given a 4-run loop, let us rescale so the loop has length 1. The run lengths are now real numbers which sum to 1, so the space of such loops is 3-dimensional. Triangle and division moves are piecewise linear transformations, where the pieces are given by linear inequalities. The set of equidistributed loops is a polyhedron, so this lemma shows that by repeated application of the triangle and division piecewise linear maps, the entire polyhedron of 4-run loops can be brought inside the equidistributed polyhedron. Figure 14 shows the equidistributed polyhedron; it is a skew-projection of a 4-dimensional cube (i.e. a *zonohedron*).

We now return to the proof. Suppose we are given a 4-run loop. Without loss of generality, let us suppose there are at least as many a's as b's, and let $G = \#a - \#b$ be the *generator inequity*, recording how many more a's than b's there are. Let x and x' be the number of b's in the longer and shorter b runs, respectively. Abusing

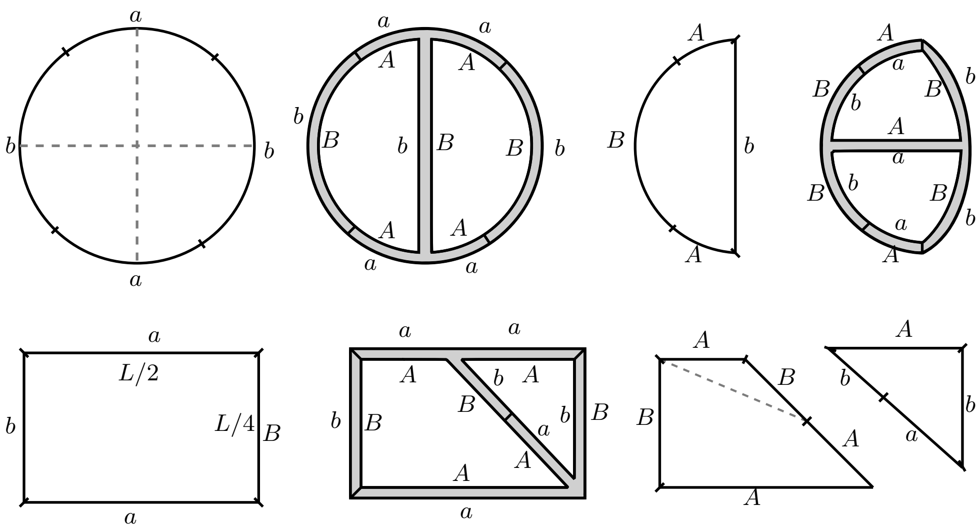

FIGURE 13. Quartering an equidistributed loop to produce exactly balanced loops and loops with a diameter between corners. Lengths are not to scale, and some lengths are labeled. The operation is read left to right, top to bottom. Notice that along the way, we handle the case of a half-equidistributed loop and a loop with adjacent runs of lengths $L/2$ and $L/4$.

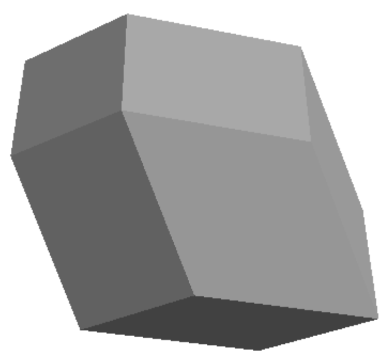

FIGURE 14. The polyhedron of equidistributed 4-run loops.

notation, we'll also refer to the runs themselves as x and x'. Note that the a or b runs may have negative exponents, so they are actually runs of A or B. For simplicity, we'll use the "positive" notation.

First, let us eliminate the degenerate case that there are very few b's. Suppose that $x + 2x' < L/2$. Consider the two diameters starting at the ends of x. There are two cases: if $x = x'$ and the runs are exactly antipodal, then both diameters run between corners, which is a case we have already dealt with. Otherwise, one of the diameters misses x' entirely, and we can cut to produce an even loop and a

loop whose generator inequity is strictly smaller (the roles of a and b are reversed, and the inequity becomes $2x' < 2(L/2 - x - x') = G$, i.e. smaller than the current inequity). See Figure 15.

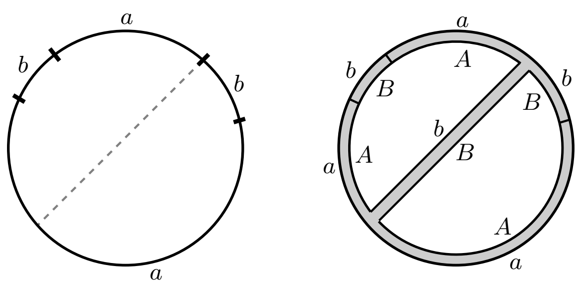

FIGURE 15. If the b runs are short, we can cut to reduce the generator inequity.

Therefore, we may assume that $x + 2x' \geq L/2$, and by definition, $x \geq x'$. We conclude that $x \geq L/6$ and $x' \geq L/4 - x/2$, and therefore that $x + x' \geq L/3$. It follows that $G \leq L/3$. Define $R = x - x'$. As a warm-up, let us first assume that $x \leq L/4$. Since $x, x' \leq L/4$, if we can find a diameter between x and x', we are done, for the diameter at right-angles will always be admissible, and the loop will be equidistributed.

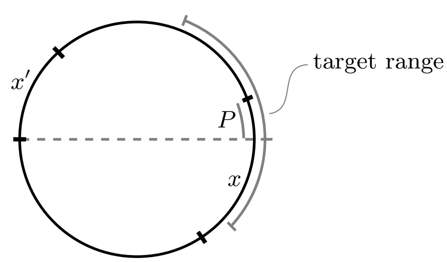

FIGURE 16. The picture showing the offset P that we'd like to minimize and the target range; we succeed in making the loop equidistributed or half-equidistributed if we shift the endpoint of x into the target range, as it is in the picture.

We always read loops counterclockwise. Let d be the diameter starting at the end of the x' run. Though we won't explicitly use it, it is useful to consider the quantity P, which is the distance between the endpoint of d and the end of x. Figure 16 shows the picture. Notice that if the endpoint of x lies anywhere in the (closed) interval indicated, then there is a diameter between x and x'. We'd like to make the magnitude of P small.

Recall that when performing a triangle move at x, we shift $L/2 - x$ of the a's from one of the a runs to the other. We compute this in terms of our variables. Note the total number of b's is $(L - G)/2$, so

$$x = \frac{\frac{L-G}{2} + R}{2} = \frac{L}{4} - \frac{G}{4} + \frac{R}{2}$$

Therefore, if we denote the number of a's which are shifted on every triangle move at x by s, we have

$$s = \frac{L}{2} - x = \frac{L}{4} + \frac{G}{4} - \frac{R}{2} \leq \frac{L}{3} - \frac{R}{2},$$

where we have used the fact that $G \leq L/3$. It is better if s is smaller, because then we can be more accurate with the placement of x. The size of the target region for the endpoint of x is $x + x'$ ($x + x' + 1$ vertices), so if $s \leq x + x'$, then the shift size is small enough that we can place the endpoint of x somewhere in the region. But we know that

$$s \leq \frac{L}{3} - \frac{R}{2} \leq \frac{L}{3} \leq x + x',$$

so the x run can be placed as desired to make the loop equidistributed.

Recall that this was the restricted case that $x \leq L/4$, which actually turns out to be the hardest. To complete the argument, we must handle the case that $x > L/4$. The issue is no longer that we can make a diameter from x' hit x, but that x might overlap the other diameter which lies at a right angle to the first. This problem is illustrated in Figure 17

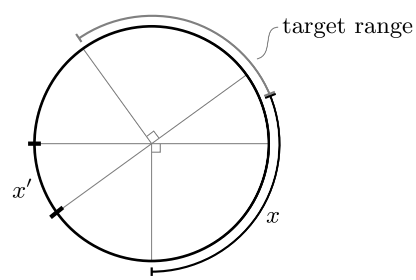

FIGURE 17. If x is large, we have to make sure it doesn't cross the diameter at a right angle to the diameter from x'.

Let us compute the size of the target range for the endpoint of x. As one can see, the target range has size $2(L/4) + x' - x = L/2 + x' - x$. But recall that the amount of a's we shift between a runs is the shift size $s = L/2 - x$, so we immediately have $s \leq L/2 + x' - x$; i.e. the shift size is smaller than the target range, so we can succeed in placing the endpoint of x in the desired range. This completes the argument showing how to make any 4-run loop equidistributed.

Since we have shown how to make any 4-run loop equidistributed, and any equidistributed loop into 2-run loops, we have reduced ourselves to a collection of loops, each of which has at most 2 runs, as desired. □

For the next step, we will modify our collection of loops with 2 runs to put them in a standard form. A *uniform* loop has a single run, so just one generator appears. An *even* loop has two runs of the same length (so length $L/2$). The *type* of a loop with two runs is a pair that records which generators appear, so for example (a, B).

Lemma 4.2.3. *Any collection of loops with* 2 *runs is pants equivalent to a collection of uniform loops, even loops, and at most one loop of each type.*

Proof. This step involves shifting and combining loops into even and uniform loops, which will leave a finite remainder. All of these operators are performed on loops

of the same type. The first step is to arbitrarily select a generator to be the *small* generator on each loop. We'll choose b. Any loop whose b run has length over $L/2$ can be cut with a diameter labeled with just a to produce an even loop and a loop with b run length less than $L/2$, as shown in Figure 18. An important feature of

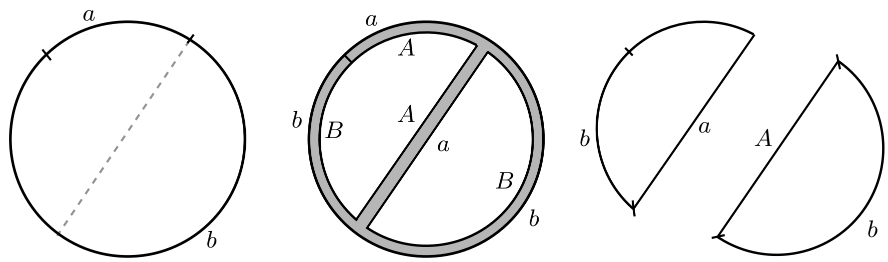

FIGURE 18. If a loop has a b run length larger than $L/2$, it can be cut to produce an even loop and a loop with a b run shorter than $L/2$. See the text for a discussion of why the inner boundary components become inverse on the right.

cutting to reduce the size of the b run is that it doesn't change the type of the loop. We have been somewhat casual about this thus far, because it makes the pictures easier to understand, but recall from the introduction to this section that if we have a loop γ, and we find a pair of pants with boundary $\gamma + \alpha + \alpha'$, then the remaining problem is to find a collection of pants with boundary $\iota(\alpha) + \iota(\alpha')$; that is, recall that γ is not pants equivalent to $\alpha + \alpha'$, but rather is pants equivalent to $\iota(\alpha) + \iota(\alpha')$. Therefore, as shown in Figure 18, the input is the loop of type (a, b), and the output is an even loop and another loop of type (a, b). The same holds true for the other operations we describe here.

We aren't concerned with the even loops, so we turn our attention to the loops with b run length strictly less than $L/2$. There are two necessary operations here; the *trade*, in which we swap pieces of the b run between loops, and the *combine*, in which we combine two small b loops into a single one (and produce a uniform loop and two even loops as byproducts). The combine operation works on any two loops whose total number of b's is strictly less than $L/2$. These operations are shown in Figures 19 and 20. Algebraically, the trade operation takes in $a^{p_1}b^{t_1+t_2}$ and $a^{p_2}b^{w_1+w_2}$, where $t_1, t_2, w_2 > 0$ and $w_1 \geq 0$, and produces $a^{p_1}b^{t_1+w_2}$ and $a^{p_2}b^{w_1+t_2}$. The combine operator takes in $a^{p_1}b^{r_1}$ and $a^{p_2}b^{r_2}$, where $r_1+r_2 < L/2$, and produces $a^{p_3}b^{r_1+r_2}$.

Notice that the trade operation requires only that *one* of the diameters be interior in the b run, or as written above, $w_1 > 0$ but $w_2 \geq 0$. Using the trade operation, we can take every b run and, for each one, move all but a single b onto a single chosen loop. Whenever this chosen loop contains a b run longer than $L/2$, we cut it off and start trading again. After this, we are left with a single loop with an unknown length b run, and possibly many loops with a single b. Then, we use the combine operation to combine these loops into a smaller number of loops with longer b runs. Then we trade the b mass to our chosen loop again, and so on.

After trading and combining as much as possible, we are left with either no loop (if the only remaining loop is even), a single loop with a b run of length less than $L/2$, or two loops whose combined run length is exactly $L/2$. This last case arises

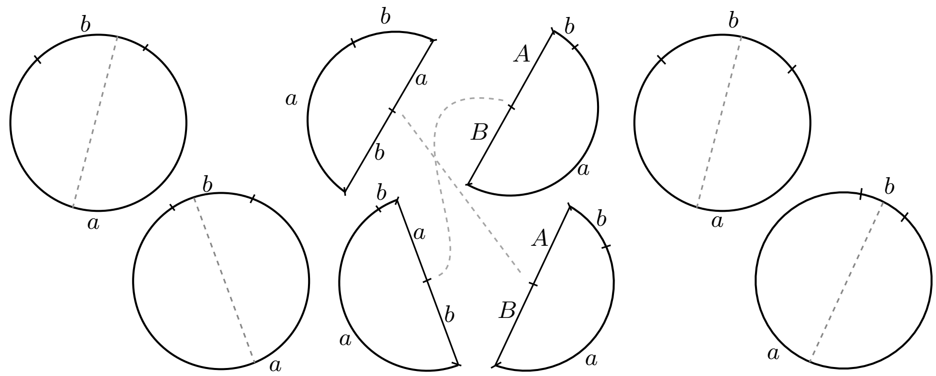

FIGURE 19. Trading pieces of b runs between loops

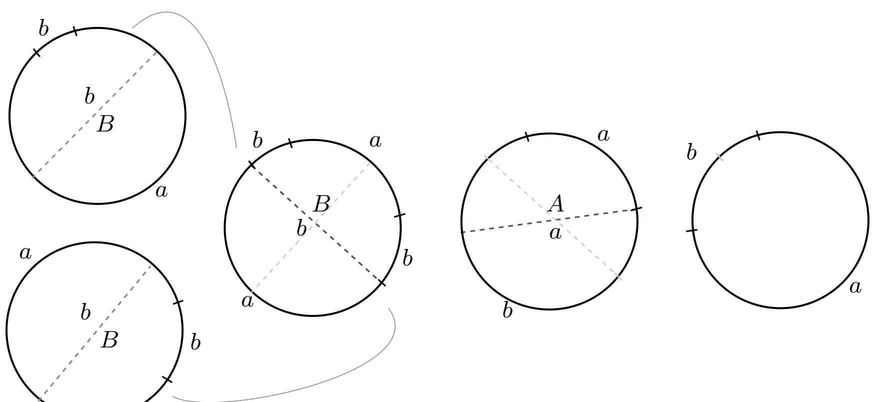

FIGURE 20. Combining b runs onto a single loop. The first step produces a byproduct uniform a loop, the second an even loop, and the third another even loop.

because we cannot use the combine operation on these loops. There is yet another sequence of moves, however, to resolve this: we use the trade operation to obtain two loops with b runs of length exactly $L/4$. Then we cut and join them to produce a perfectly balanced loop with 4 runs. A single triangle move results in an even loop plus a commutator; we then apply Figure 11 to the commutator to get a pair of inverse loops.

Doing this to each loop type proves the lemma. □

The final step in the proof of Proposition 4.1.1 is to show that we can attach pants and annuli to the output of Lemma 4.2.3, that is, uniform loops, even loops, and a single remainder loop of each type, so that we have nothing left. Let $x_{a,b}$, $x_{a,B}$, $x_{A,b}$, and $x_{A,B}$ denote the run length of the single b run in the remainder loop of each type. Rescaling and considering arbitrary L, we can think of each variable as a real number in the open interval $(0, 1/2)$. The fact that the entire collection

of loops must be homologically trivial gives us two linear equations

$$x_{a,b} + x_{a,B} - x_{A,b} - x_{A,B} = k_1$$
$$(1 - x_{a,b}) - (1 - x_{a,B}) + (1 - x_{A,b}) - (1 - x_{A,B}) = k_2,$$

Since there are even and uniform loops to consider, it is not *a priori* the case that $k_1 = k_2 = 0$. However, it *is* the case that $k_1, k_2 \in \frac{1}{2}\mathbb{Z}$, and $k_1 - k_1 \in \mathbb{Z}$, since the uniform and even loops change homology discretely. Solving the linear equations shows that the only solutions to this linear system in $(0, 1/2)$ exist when $k_1 = k_2 = 0$, and in fact it must be that $x_{a,b} = x_{A,B}$ and $x_{a,B} = x_{A,b}$. Since these variables are equal, we can use annuli to glue the remainder loops. The homologically trivial collection of uniform and even loops can now be glued along entire edges, so is obviously pants equivalent to the empty collection. That is, we have successfully produced a collection of pants and annuli with boundary $s' + t + \iota(t)$, where s' is our original collection, and t is the many intermediate boundaries we used to reduce s' to the empty set. Adding in ι-annuli, then, gives us a collection of pants and annuli which has boundary $s' + n\mathbf{1}'$, for some sufficiently large n.

Observe that we never need to duplicate our collection s', or, equivalently, multiply v' by any factor. We have therefore proved the stronger Proposition 4.1.2 for rank 2 free groups.

4.3. **Higher rank.** We have completed the proof of Proposition 4.1.1 in the case that the free group has rank 2. We now describe the necessary modifications to the argument for higher rank free groups. Given a collection of loops $s' \in S'(L)$, the same technique of cutting with diameters works to show that s' is pants equivalent to a collection of 4-run loops. However, triangle moves no longer apply, since each of the 4 runs might be a run of a different generator.

The first step is to attach pants in such a way that we are left with 4-run loops, each of which only involves two generators. This is actually quite straightforward, since we have more freedom with the labels on the diameters that we attach. Figure 21 shows how to attach diameters to produce loops of the desired form which are pants equivalent to the original loop. Technically, the figures represent simply unions of pants, not the (non-trivalent) fatgraphs shown. They are drawn as shown to emphasize the point that we produce several other byproduct loops, but they come in cancelling inverse pairs. All the interior diameters shown have length $L/2$, even though they are not drawn to scale.

We remark that the pictures in Figure 21 are general, up to rotation and reflection. The double-diameter from top to bottom exists by the argument in the proof of Lemma 4.2.1, and the double-diameter from the lower left corner to the middle exists because the first diameter has length $L/2$. We also remark that in higher rank, it is possible that we have some 3-run loops. It is simple to cut these to 4-run loops and then apply the above argument.

After applying Lemma 4.2.2, we may assume that we are left entirely with uniform loops, even loops, and one 2-run loop of each type. Now, though, there are $4\binom{r}{2}$ loop types, which is too many to duplicate the linear-algebraic argument from the previous section. The simple solution is to take $L/2$ copies of our collection. Now each loop type is repeated exactly $L/2$ times, so when we re-collect the remainder, we are left with no remainder, so we have only uniform loops and even loops, which can be paired arbitrarily. Therefore, we have found a collection of

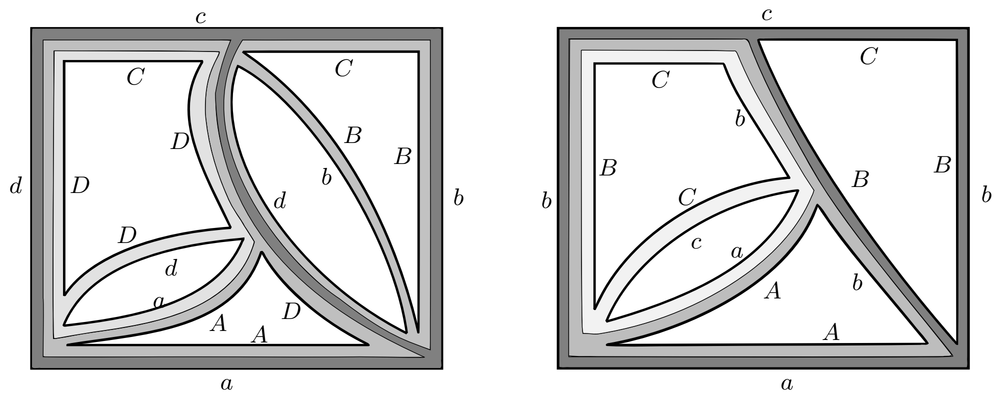

FIGURE 21. Reducing each loop to lie in a rank 2 subgroup. The fatgraph on the left is built out of folded pants only when b and d are distinct generators. If $b = d^{\pm 1}$, then we use the picture on the right. A similar picture holds when the vertical diameter has endpoints on runs of the same generator. The point is that all of the non inverse-matched pants boundaries lie in rank 2 subgroups.

pants and annuli which has boundary $(L/2)v' + n\mathbf{1}'$, which completes the proof of Proposition 4.1.1.

5. Random one-relator groups

Throughout this section we fix a free group F_k with $k \geq 2$ generators, and we fix a free generating set. For some big (unspecified) constant n, we let r be a random cyclically reduced word of length n, and we consider the one-relator group $G := \langle F \mid r \rangle$.

In this section we will show that with probability going to 1 as $n \to \infty$, the group G contains a surface subgroup $\pi_1(S)$. The surface S in question can be built from N disks bounded by r and N disks bounded by r^{-1}, glued up along their boundary in such a way that the 1-skeleton is a trivalent fatgraph with every edge of length $\geq L$, where L is some (arbitrarily big) constant fixed in advance, and $N \leq 20L$ is the constant in the Thin Fatgraph Theorem 3.3.1.

The group G is evidently $C'(\lambda)$ for any $\lambda > 0$ with probability going to 1 as $n \to \infty$, and therefore the injectivity of S can be verified by showing that the 1-skeleton of S does not contain a long piece in common with r or r^{-1}, except for a piece contained in the boundary of one of the $2N$ disks.

A trivalent graph Y in which every edge has length at least L has at most $|Y| \cdot 2^{m/L}$ subpaths of length m, where $|Y|$ is the length of Y. In a free group of rank k there are (approximately) $(2k-1)^m$ reduced words of length m, and the relator r contains at most n of them. Taking $m = \epsilon \cdot n$, a simple counting argument shows that Y does not have any path of length m in common with an *independent* random word of length n. However, Y and r are utterly dependent, and we must work harder to show that S is injective.

The key idea is to break up $Nr \cup Nr^{-1}$ into pieces (called *beads*) which each bound their own trivalent fatgraph. Then subpaths in the fatgraph associated to one bead *will* be independent of the subpaths of $Nr \cup Nr^{-1}$ associated to another bead, and this argument can be made to work.

5.1. **Independence and correlation.** A random word in a finite alphabet (in the uniform distribution) has the property that any two disjoint subwords are independent. A random *reduced* word in the free group fails to have this property, since (for example) if uv are adjacent subwords, the last letter of u must not cancel the first letter of v (so the words are not really independent). However, such words have a slightly weaker property which is just as useful as independence in most circumstances; this property can be summarized by saying that *correlations decay exponentially*. This means that if u and v are reduced words, and we fix some distance T, then the probability of finding some substring v at a specific location *conditioned* on finding u at a location a distance T from it, differs from $1/(2k)(2k-1)^{v-1}$ (i.e. the probability if disjoint subwords were really independent) by $(2k-2)^{-T}$. See e.g. [3], Lem. 2.4 for a careful proof.

If $T \gg |v|$ (for instance, if $|v| = O(\log(n))$ while $T = O(n^\epsilon)$) then in practice one can treat u and v as though they were independent. In the sequel we will say informally that *correlations* (between disjoint subwords) *decay exponentially.*

5.2. **Beads.** Let r be a random cyclically reduced word of length n. We fix some (small) positive constants C and δ; later we will say how small they should be.

Definition 5.2.1. A *bead decomposition* of r is a decomposition of r into segments labeled (in cyclic order) $r_0, r_1^+, r_2^+, \cdots, r_{M-1}^+, r_M, r_{M-1}^-, \cdots, r_1^-$ where each $r_i^\pm$ has length $n^{1-\delta} \cdot (1/2 + o(1))$ and r_0, r_M have length $n^{1-\delta} \cdot (1 + o(1))$ (so that $M = n^\delta \cdot (1 + o(1))$) so that the prefix of r_i^+ of length $C\log(n)$ is inverse to the suffix of r_i^-, and the prefix of r_M of length $C\log(n)$ is inverse to the suffix of r_M.

Given a bead decomposition of r, we glue the mutually inverse subwords described above, thereby decomposing r into a sequence of loops ("beads") of length $n^{1-\delta} \cdot (1 + o(1))$ with one or two tags. Denote these beads $B_0, B_1, \cdots, B_M$. The edges of length $C\log(n)$ obtained by gluing the inverse prefixes/suffixes we refer to as the *lips* of the beads.

Lemma 5.2.2 (Bead Decomposition)**.** *There exists a bead decomposition with probability $1 - O(e^{-n^c})$. Moreover, the bead decomposition can be chosen so that the correlations between distinct B_i decay exponentially.*

Proof. We start at an arbitrary location in r, and take this to be the approximate midpoint of r_0. Extend this region $n^{1-\delta} \cdot (1/2 - n^{-\delta})$ in either direction, and then read the next pair of segments of length $n^{1-2\delta}$ synchronously until the first time we read off a pair of mutually inverse segments of length $C\log(n)$. If $C < (1-2\delta)/\log(2k-1)$ then such a pair of mutually inverse segments will occur with probability $1 - O(e^{-n^c})$ by Chernoff's inequality, for some $c > 0$ depending on δ. We then build successive segments $r_i^\pm$ in order by the same procedure. There are n^δ such segments, and this polynomial term can be absorbed into the exponential estimate of probability at the cost of adjusting constants. To see that correlations between different B_i decay exponentially, imagine that we generate the word r by a Markov process letter by letter as we go, so that we only generate successive B_i after having already constructed B_j with $j < i$. □

There is no reason to expect that the beads B_i are homologically trivial, but there is a trick to adjust them so that they are. We build a bead decomposition of r and of r^{-1} simultaneously, so that the beads of r^{-1} have inverse labels to the

beads of r. We denote the beads of r^{-1} by B_i^{-1}, and note that they are inverse (as tagged cyclic words) to the B_i. Then for each i the union $B_i \cup B_i^{-1}$ is homologically trivial.

By Theorem 3.3.1 for each i, and for some $N \le 20L$, the collection $NB_i \cup NB_i^{-1}$ bounds a trivalent fatgraph Y_i with all edges of length at least L, with probability $1 - O(e^{-n^c})$. By Lemma 5.2.2, for $j \ne i$ the beads B_j are almost uncorrelated with the Y_i, so for any ϵ there is $c(\epsilon)$ so that with probability $1 - O(e^{-n^c})$ there are no paths in Y_i of length n^ϵ in common with any segment in $r - r_i^{\pm}$.

Lemma 5.2.3 (No long piece)**.** *Let Y be the fatgraph obtained from the union of the Y_i associated to a bead decomposition as above. Then with probability $1 - O(e^{-n^c})$ every path in Y of length $\epsilon \cdot n$ which appears in r or r^{-1} is in $\partial S(Y)$.*

Proof. Let γ be a path in Y of length $\epsilon \cdot n$ and let $\gamma' \subset r$ (without loss of generality) have the same labels as γ. Then for any $\epsilon > 0$ there is a $c > 0$ so that for each i and each subsegment σ' of γ' of length n^ϵ contained in B_i the corresponding subsegment σ' of γ' must have at least $(1-\epsilon)$ of its length contained in Y_i. By the definition of the bead decomposition, successive subsegments of γ' in adjacent B_i are joined by paths of length $C\log(n)$ running over the lip. The corresponding subsegments in γ that transition from Y_i to Y_{i+1} must *also* run over the lip, so there is another copy of the word on the lip within n^ϵ of the lip. If the two copies are not distinct, so that γ and γ' overlap on a common piece, then since Y is folded we must simply have $\gamma = \gamma'$ and γ is in $\partial S(Y)$ as claimed. Otherwise there are two distinct copies of the lip contained in a segment of length n^ϵ in r. If ϵ is sufficiently small compared to C, the probability that two identical subwords of length $C\log(n)$ will occur in a specific segment of length n^ϵ is arbitrarily small (in fact, of size $O(n^{-\epsilon'})$ for some ϵ' depending on ϵ and C). The probability of such occurrences at adjacent lips is not independent, but the correlations decay exponentially, so since this must happen for $O(n^\delta)$ successive segments Y_i (the segments that γ' intersects in order), the probability of all these lips having nearby distinct copies is $O(e^{-n^c})$ and the lemma is proved. $\square$

We deduce the following theorem as a corollary:

Theorem 5.2.4 (Random One-Relator)**.** *Fix a free group F_k and let r be a random cyclically reduced word of length n. Then $G := \langle F_k \mid r \rangle$ contains a quasiconvex surface subgroup with probability $1 - O(e^{-n^c})$.*

Proof. This follows from Lemma 5.2.3 and Lemma 2.2.5. $\square$

In fact, there are $O(e^{n^c})$ many choices of bead decomposition, and most of these give rise to quasiconvex surface subgroups.

Definition 5.2.5. We call the surfaces constructed as above *beaded surfaces*.

Note that a beaded surface has genus $O(n)$. In fact, in the case that $k = 2$, a beaded surface can actually be taken to have genus $o(n)$, since we can take $N = 1$ in the application of the Thin Fatgraph Theorem, and then as $n \to \infty$ we can take $L \to \infty$. It seems very likely that the surfaces can be taken to have genus $o(n)$ for any k.

6. Random groups

In this section we prove our main theorem, that a random group at density $D < 1/2$ contains a surface subgroup with probability $1 - O(e^{-n^c})$. In fact, our argument shows that it contains *many* subgroups (of genus $O(n)$). Our argument depends on some elements of the theory of small cancellation developed for random groups by Ollivier [12], and we refer to that paper several times.

6.1. **Small cancellation in random groups.**

For later convenience, we here state three results from Ollivier [12] that we use in the sequel.

Theorem 6.1.1 (Ollivier, [12], Thm. 2)**.** *Let G be a random group at density D. Then for any positive ϵ, and any reduced van Kampen diagram $\mathcal{D}$ containing m disks, we have*

$$|\partial\mathcal{D}| \geq (1 - 2D - \epsilon) \cdot nm$$

with probability $1 - O(e^{-n^c})$

Here the hardest part is to show that the same ϵ works for van Kampen diagrams of *arbitrary* size.

Theorem 6.1.2 (Ollivier, [12], Cor. 3)**.** *Let G be a random group at density D. Then the hyperbolicity constant δ of the presentation satisfies $\delta \leq 4n/(1-2D)$ with probability $1 - O(e^{-n^c})$.*

Theorem 6.1.3 (Ollivier, [12], Thm. 6)**.** *Let G be a random group at density D. Then for any positive ϵ, and for any reduced van Kampen diagram $\mathcal{D}$ with at least two faces, there are at least two faces which have a (connected) piece on $\partial\mathcal{D}$ of length at least $n(1 - 5D/2 - \epsilon)$, with probability $1 - O(e^{-n^c})$.*

The statements of theorems in Ollivier's paper do not make the estimate of probability (as a function of n) explicit; however these estimates are straightforward to derive from his methods (and in any case, we do not use them in the sequel).

6.2. **Convexity and (m, α)-convexity.**

Definition 6.2.1. Fix a free group F_k with $k \geq 2$ and a free generating set. A *random group of length n and density $D < 1$* is obtained from F_k by adding $(2k-1)^{nD}$ independent random cyclically reduced relations of length n.

Gromov showed that a random group at density D satisfies the small cancellation condition $C'(2D)$ with probability going to 1 as $n \to \infty$. Pick one relator r, and build a beaded surface S as in § 5 whose spine is trivalent and with every edge of length $\geq L$ for some large (fixed) L.

Lemma 6.2.2. *Fix D. Then for any $\alpha > D$, a beaded surface S constructed by the method of § 5 is α-convex, with probability $1 - O(e^{-n^c})$.*

Proof. By Lemma 5.2.3 for any positive ϵ, the spine Y of S has no piece of length $\epsilon \cdot n$ in common with the relator r except for pieces occurring in $\partial S(Y)$. For any $\alpha > 0$ there are $O(2^{\alpha n/L})$ paths in Y of length αn. Define $\delta = \log(2)\alpha/L$ so that $e^{\delta n} = 2^{\alpha n/L}$, and note that by taking L sufficiently large, we can make δ as small as we want.

There are $(2k-1)^{\alpha n}$ reduced words of length αn, and a random relator r' contains n subwords of this length, so a random relator r' has probability $O(e^{\delta n - \log(2k-1)\alpha n})$

of having a piece of length αn in common with Y. If $\alpha > D$ and $\delta < \log(2k-1)(\alpha - D)$ then no relator $r' \neq r$ has a subpath of length αn in common with Y, with probability $O(e^{-n^c})$. $\square$

We deduce by Lemma 2.2.5 that a random group contains a surface subgroup at any $D < 1/12$. However, Theorem 6.1.3 already improves this to $D < 2/7$.

To go further we need two ingredients — we need to generalize the condition of α-convexity to rule out the existence of relators with several short pieces in common with Y, and we need to control the dependence of different faces in a big diagram.

Definition 6.2.3 ((m, α) convexity)**.** Let G be a group with a fixed presentation, and let S be an oriented surface over the presentation with a folded spine Y. We say S is (m, α)-*convex* (for some integer $m > 0$ and some real $\alpha > 0$) with respect to the presentation if for every immersed path γ in Y which contains m disjoint subpaths γ_j which are disjoint pieces in some relation $r_i^{\pm}$ with $|\gamma_j|/|r_i| \geq \alpha$, we actually have that γ is contained in $\partial S(Y)$ (i.e. it is in the boundary of a disk of S).

Note that an (m, α)-convex surface is $m\alpha$-convex, but the converse is not necessarily true.

Lemma 6.2.4. *Fix D. Then for any m, α with $m\alpha > D$, a beaded surface S is (m, α)-convex with probability $1 - O(e^{-n^c})$.*

Proof. The proof is essentially the same as that of Lemma 6.2.2. The probability that a random relator r' has a piece of length αn in common with the spine Y is $O(e^{\delta n - \log(2k-1)\alpha n})$, and the probability of having m disjoint pieces of that length is therefore $O(e^{m(\delta n - \log(2k-1)\alpha n)})$. So for $m\alpha > D$ the result follows as before. $\square$

6.3. **van Kampen disks.** Our strategy will be to show that the existence of a certain kind of van Kampen disk $\mathcal{D}$ with boundary a cyclically reduced word in Y essential in S gives rise to a contradiction. Suppose that $\gamma \subset Y$ is an essential loop in S whose image is trivial in G, so that there is some van Kampen diagram $\mathcal{D}$ with boundary γ. If some face in $\mathcal{D}$ has boundary r or r^{-1}, and if this face has more than ϵn in common with γ, then this face agrees with a disk of S, and we can find a smaller van Kampen diagram $\mathcal{D}'$ by pushing across this disk. So in the sequel we will only consider loops $\gamma \subset Y$ essential in S bounding van Kampen disks $\mathcal{D}$ which cannot be simplified by such a move. We call such a van Kampen disk *efficient*.

The following Lemma is standard.

Lemma 6.3.1 (Short shortcut)**.** *Let G be a hyperbolic group with a presentation with respect to which it is δ-hyperbolic. Let Γ be a cyclically reduced word in the generators which is trivial in G. Then there is a van Kampen disk $\mathcal{D}$ with $|\partial\mathcal{D}| \leq 18\delta$ and a connected subpath $\gamma \subset \partial\mathcal{D}$ with $\gamma \subset \Gamma$ and $|\gamma| > |\partial\mathcal{D}|/2$.*

Note that if $\gamma' = \partial\mathcal{D} - \gamma$ then $|\gamma'| < |\gamma|$. In other words, γ' is a *shortcut*; hence the terminology.

Proof. In any δ-hyperbolic path metric space, for any $k > 8\delta$, a k-local geodesic (i.e. a 1-manifold for which every subpath of length at most k is a geodesic) is a (global) $(\frac{k+4\delta}{k-4\delta}, 2\delta)$-quasigeodesic; see [2], Ch. III. H, 1.13 p. 405. The loop Γ starts and ends at the same point, and is therefore not a k-local geodesic for $k \geq 9\delta$. Therefore some segment of length at most 9δ is not geodesic, and it cobounds a van Kampen disk $\mathcal{D}$ with an honest geodesic. $\square$

We deduce the following corollary:

Lemma 6.3.2. *Suppose that S is a beaded surface which is not π_1-injective. Then there are constants C and C' depending only on $D < 1/2$, a path γ in the spine Y (not contained in $\partial S(Y)$) of length at most Cn, and a van Kampen diagram $\mathcal{D}$ containing at most C' faces so that $\gamma \subset \partial\mathcal{D}$ and $|\gamma| > |\partial\mathcal{D}|/2$.*

Proof. Theorem 6.1.2 says that $\delta \leq 4n/(1-2D)$, so by Lemma 6.3.1 it follows that there is such a disk $\mathcal{D}$ with boundary of length at most $72n/(1-2D)$. On the other hand, by Theorem 6.1.1 we know $72n/(1-2D) \geq |\partial\mathcal{D}| \geq (1-2D-\epsilon)\cdot nC'$ where C' is the number of faces; in particular, C' is *bounded* in terms of D (and *independent* of n). $\square$

The fact that C and C' can be chosen *independent* of n (but depending on $D < 1/2$ of course) is crucial for our purposes.

6.4. **Surfaces in random groups.** We can now prove the main theorem of the paper.

Theorem 6.4.1 (Surfaces in random groups)**.** *A random group of length n and density $D < 1/2$ contains a surface subgroup with probability $1-O(e^{-n^c})$. In fact, it contains $O(e^{n^c})$ surfaces of genus $O(n)$. Moreover, these surfaces are quasiconvex.*

Proof. Pick one relation r and build a beaded surface S by the method of § 5. We have already shown that the 1-skeleton Y of S does not contain a path of length $\epsilon\cdot n$ in common with r or r^{-1} except for paths in the boundary of a disk, with the desired probability. By Lemma 6.2.4 the surface is (m,α)-convex for any $1/2 > m\alpha > D$.

Suppose S is not π_1-injective. Then by Lemma 6.3.2 there is an efficient van Kampen disk with boundary an essential loop in Y, containing a subdisk $\mathcal{D}$ with at most C' faces, and at least half of its boundary equal to some path γ in the spine Y. Since the original disk was efficient, no face of $\mathcal{D}$ with boundary label r or r^{-1} has more than ϵn in common with γ.

Fix a combinatorial type for the diagram. Then there are at most polynomial in n choices of edge lengths for the edges in the diagram. Choose a collection of edge lengths. Let $m < C'$ be the number of faces.

We estimate the probability that there is a way to label each face with a relator or its inverse compatible with γ. We express the count in terms of *degrees of freedom*, measured multiplicatively, as powers of $(2k-1)$ (Gromov uses the terms *density* and *codensity*; see the discussion in [8] pp. 269–272 expanded at length in [13]).

Each choice of face gives nD degrees of freedom, and each segment in the interior of length ℓ imposes ℓ degrees of constraint. Similarly, the segment γ imposes $|\gamma|$ degrees of constraint. Let $\mathcal{I}$ denote the union of interior edges. Then $|\mathcal{D}|+2|\mathcal{I}| = nm$ so $|\gamma|+|\mathcal{I}| = nm/2$ because $|\gamma| \geq |\partial\mathcal{D}|/2$. On the other hand, the total degrees of freedom is $nmD < nm/2$, so no assignment is possible, with probability $1-O(e^{-n^c})$. Summing the exceptional cases over the polynomial in n assignments of lengths, and the *finite* number of combinatorial diagrams, we see that S is injective, with probability $1-O(e^{-n^c})$ for some c depending on D (and going to 0 as $D \to 1/2$).

In fact, the same argument implies that every essential path in Y with no long segment on $\partial S(Y)$ is actually quasigeodesic, by the proof of Lemma 6.3.1. The theorem follows. $\square$

Remark 6.4.2. The surfaces we build are homologically trivial, since they map nontrivially over only one disk bounded by a relator r, and with total degree 0. It turns out that essentially the same method produces many homologically essential surfaces.

If n is even, a random reduced word of length n in a free group of rank k is homologically trivial with probability $O(n^{-k/2})$. Since there are $(2k-1)^{nD}$ relators, there are an enormous number of such homologically trivial relators, and we can try to build a surface mapping over the associated disk with degree 1 (and therefore being homologically essential in G). Evidently, the only obstruction to finding such surfaces is to build a bead decomposition as in Lemma 5.2.2 where all the B_i are homologically trivial, while still preserving the property that correlations between distinct B_i decay exponentially fast. The probability that the naive construction of a bead decomposition (as in the lemma) applied to a random word will have this property is $(n^{-k/2})^{n^\delta}$, which is subexponential in n, so many of the relators will have this property, and we can build many homologically essential surfaces of genus $O(n)$. If n is odd we can build a similar (homologically essential) beaded surface from two (judiciously chosen) relators.

Remark 6.4.3. The surfaces we build have genus $O(n)$ (or $o(n)$ for rank 2), and it is natural to wonder if this is the best possible. We conjecture not; in fact we conjecture that the smallest genus injective surfaces in random groups are of genus $O(n/\log n)$ (at any density $D < 1/2$).

In fact, Thm. 4.16 of [3] lets us give a precise estimate of the geometry of the *Gromov norm* on $H_2(G;\mathbb{R})$. Let V be the vector space with the relators r_i as basis, and let W be the kernel of the natural map $V \to H_1(F_k;\mathbb{R})$. Then we can identify W with $H_2(G;\mathbb{R})$, by Mayer–Vietoris. The vector space W inherits an L^1 norm from V with respect to its given basis. The Gromov norm on W (on random subspaces of fixed dimension) is (with overwhelming probability) proportional to this L^1 norm, with constant of proportionality $2\log(2k-1)n/3\log(n)$, up to a multiplicative error of size $1+o(1)$ (there is a factor of 4 relative to the statement of Thm. 4.16 of [3]; this factor of 4 reflects the difference between the Gromov norm and the so-called *scl norm*). It seems plausible that classes in $H_2(G;\mathbb{Q})$ should be projectively represented by norm-minimizing surfaces; such surfaces will necessarily be injective. Again, it seems likely that one should not need to pass to a very big multiple of a class to find an extremal surface (at least for some classes); so there should be injective surfaces of genus $O(n/\log(n))$. We strongly suspect this order of magnitude is sharp.

References

[1] I. Agol (with an appendix with D. Groves and J. Manning), *The virtual Haken conjecture*, arXiv:1204.2810

[2] M. Bridson and A. Haefliger, *Metric spaces of nonpositive curvature*, Grund. der math. Wiss. Springer-Verlag Berlin 1999

[3] D. Calegari and A. Walker, *Random rigidity in the free group*, Geom. Topol. to appear

[4] D. Calegari and A. Walker, `scallop`, Computer program available from the authors' webpages.

[5] D. Calegari and A. Walker, *Surface subgroups from linear programming*, arXiv:1212.2618

[6] D. Calegari and H. Wilton, *Random graphs of free groups contain surface subgroups*, arXiv:1303.2700

[7] F. Dahmani, V. Guirardel and P. Przytycki, *Random groups do not split*, Math. Ann. **349** (2011), no. 3, 657–673

[8] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, vol. 2 LMS Lecture Notes **182** (Niblo and Roller eds.) Cambridge University Press, Cambridge, 1993
[9] M. Gromov, personal communication
[10] J. Kahn and V. Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. Math. **175** (2012), no. 3, 1127–1190
[11] M. Kotowski and M. Kotowski, *Random groups and Property (T): Zuk's theorem revisited*, arXiv:1106.2242
[12] Y. Ollivier, *Some small cancellation properties of random groups*, Internat. J. Algebra Comput. **17** (2007), no. 1, 37–51.
[13] Y. Ollivier, *A January 2005 invitation to random groups*, Soc. Bras. de Mat. Ens. Mat. **10**, 2005
[14] Y. Ollivier and D. Wise, *Cubulating groups at density* $< 1/6$, Trans. AMS **363** (2011), no. 9, 4701–4733
[15] J. Stallings. *Topology of finite graphs*, Invent. Math. **71** (1983), no. 3, 551–565
[16] A. Zuk, *Property (T) and Kazhdan constants for discrete groups*, Geom. Func. Anal. **13** (2003), no. 3, 643–670

Department of Mathematics, University of Chicago, Chicago, Illinois, 60637
E-mail address: `dannyc@math.uchicago.edu`

Department of Mathematics, University of Chicago, Chicago, Illinois, 60637
E-mail address: `akwalker@math.uchicago.edu`